

All networks look the same to me: Testing for homogeneity in networks

Jono Tuke and Matthew Roughan

December 4, 2015

Abstract

How can researchers test for heterogeneity in the local structure of a network? In this paper, we present a framework that utilizes random sampling to give subgraphs which are then used in a goodness of fit test to test for heterogeneity. We illustrate how to use the goodness of fit test for an analytically derived distribution as well as an empirical distribution. To demonstrate our framework, we consider the simple case of testing for edge probability heterogeneity. We examine the significance level, power and computation time for this case with appropriate examples. Finally we outline how to apply our framework to other heterogeneity problems.

1 Introduction

There are many examples of complex-interaction systems, often described as *networks*, for which we can have only a single example: *e.g.*,

- the phylogenetic tree describing the evolution of species [Huelsenbeck and Ronquist, 2001];
- the Internet [Roughan et al., 2011];
- the global inter-species food web [Dunne et al., 2002];
- the scientific collaboration network [Newman, 2001]; or
- the complete human social network [Wasserman and Faust, 1994].

Although these networks are sometimes considered as pluralities (as the references above often do), there is actually a single network from which we observe smaller components. The network might evolve over time, but snapshots of its evolution are highly dependent samples of the network in question. So in reality, we have one sample of each.

There is a philosophical problem in modelling a system for which we have a single data point, let alone a high-dimensional system with only one datum. Namely, how can we balance the conflicting demands in modelling: on the one hand we want a model

that correctly fits the observations; and on the other hand a model that is simple and explanatory. The former extreme is represented by “the data is the model” (which holds zero explanatory power), and the later by the model that all datasets are the same, and any discrepancy is just noise (a ludicrous supposition, presented only as a example of the extreme).

Somewhere in between lies what Box and Draper would call a “useful” model [Box and Draper, 2007], but how to know? Statistics has been employed for more than 200 years towards answering that very question. For instance, a model might commonly be tested through some form of cross-validation where some data-points are removed from the fitting process in order to provide a test set. If we have only one datum this is clearly impossible, as are all simple statistical tests.

Consequently, the analysis and modelling of such systems can take one of two forms¹:

1. We know the “physics” of the system, at least to some approximation, and exploit this side-information in our analysis. For instance, knowing that a network was generated by a process of growth with preferential attachment [Barabási and Albert, 1999], we could use the data simply to estimate the growth parameters².
2. Or we make the critical assumption that smaller components of the model can be used to obtain samples from which we can draw statistics, and thereby make an assessment of the correct model. For instance, in the analysis of food webs, we might examine (approximately) isolated groups, *e.g.*, Dunne et al. [2002].

Of the two methods, the second approach is valued for its explanatory power, *i.e.*, from it we can derive new physics. However, the critical assumption must be valid. In the analysis of subpopulations of species the assumption may well be, as it might in a network that can decompose into only loosely coupled components, each of which can be used to provide pseudo-independent samples of the larger network. But sometimes, the assumption is worrying.

The assumption requires that local structure of a network is homogenous. In our context that means that each subgraph is statistically the same. If not, we have the possibility of what is sometimes called a Type-III error “giving the right answer to the wrong problem” [Kimball, 1957]. That is, a perfectly reasonable statistical procedure may provide results that would be correct given the assumption, but are actually far from reasonable.

So how can we test the relationship between the global structure of a network and the local structure of its subgraphs? Let’s illustrate with a simple example of the Gilbert-Erdős-Rényi (GER) network. In GER networks, the probability of an edge between any two nodes is constant (See Section 2). What if we have a network where the probability of

¹Of course there are other approaches, but they often involve flawed logic, *i.e.*, model X has feature A, and we observe A, therefore model X is correct.

²Note that using the data for parameter estimation from a model is very different from using the data to choose the model!

an edge depends on a property of the nodes. Now the subgraphs will have heterogeneity in the overall probability of edges.

How do we test for this heterogeneity, how do examine the relationship between the local structure of subgraphs and the global structure of a network?

In this paper, we illustrate how this is achieved using sampling from the network. This paper is a proof of concept with a simple example to illustrate its potential. We consider two simple tests to test for homogeneity of edge probability. These methods use sampling without replacement and goodness of fit tests. We outline the algorithm for both methods and test their significance, power, and computational cost and show that they can provide practical analysis of random networks.

Ultimately, the goal is to provide a path towards diagnostic tests of network models. That is, rather than approaching modelling as a craft, requiring expertise and deep knowledge of graph theory, we should be able to provide a standard suite of estimators, along with diagnostics for those estimators that can be used by any practitioner, much as modern statistics has for regression. For instance, on performing linear regression, one might then test for heteroscedasity. Here we propose fitting random-graph models, but with formal means to test underlying assumptions inherent in the test.

The exemplars we present are simple, but the advantage of the underlying idea (as opposed to a test designed specifically for a particular model) is that it is easily generalised; it is simply a matter of choosing an appropriate sampling, and variate against which to measure homogeneity. Then one can consider and test for more general notions of heterogeneity.

2 Graphs and Sampling

Consider an undirected network (V, E) such that the (i, j) entry of the adjacency matrix is denoted A_{ij} and is defined as

$$A_{ij} = \begin{cases} 1, & \text{edge between node } i \text{ and node } j, \\ 0, & \text{no edge between node } i \text{ and node } j. \end{cases}$$

The networks we consider are undirected with no loops so $A_{ij} = A_{ji}$, and $A_{ii} = 0$, but the results are generalizable. Denote the number of edges and nodes as $|E|$ and $|V|$, respectively.

We denote by p_{ij} the probability $P\{A_{ij} = 1\}$, and make the following definition:

Definition 1 *We call a network homogeneous with respect to edge probability if the edge probability $p_{ij} = p$ is constant, i.e., the edge probability does not depend on any (potentially hidden) node properties.*

Note two features of the definition:

1. We define homogeneity with respect to a feature of the graph. A common source of confusion in the term seems to be that different authors use homogeneity with respect to alternative features: here we suggest that the definition must be explicit, but the definition generalizes in the sense that we could easily incorporate other variates in place of edge probability.
2. The definition here is equivalent to that of GER random graphs $G(n, p)$, but that is not the aim. Our goal is to test for a *feature* of the data, not a specific model.

The latter is an important point in general. Much statistical modelling is about fitting models or estimating parameters, given certain assumptions. This paper is aimed at testing assumptions. The distinction is important: as a result we will not present the common approach to test for a GER graph by examining the node degree distribution, as calculation of this distribution inherently presumes *a priori* that the network is homogeneous, and that we can estimate the distribution by examining the statistics of all nodes as if they were identically distributed samples drawn from an underlying variate.

We present two methods for testing if an observed network is homogeneous with respect to edge probability, *i.e.*,

$$H_0 : p_{ij} = p$$

$$H_a : \text{at least one } p_{ij} \neq p.$$

Note, however, that we only know the A_{ij} , not the p_{ij} . We can't even form reasonable estimates of the individual p_{ij} , as we have a single sample of each.

Both methods we present use a goodness of fit test applied to the number of edges observed in sampled subgraphs from the observed network. To obtain a sampled subgraph, $G'(V', E')$ we sample $k = |V'|$ nodes (without replacement) from the observed network $G(V, E)$ and choose E' such that

$$(i, j) \in E', \text{ iff } i, j \in V' \text{ and } (i, j) \in E.$$

The strategy above is referred to as *node sampling* (see Lee et al. [2006]). However, once again, note that the sampling procedure is arbitrary. We have chosen one of the simplest possible here, but there are alternative approaches such as link and snowball sampling Lee et al. [2006].

3 Testing for Homogeneity

The statistics of sampled subgraphs have been studied, however, that study seems to have focussed on standard graph metrics, such as power-law degree exponents, *e.g.*, see Lee et al. [2006]. Here we are interested in statistics that are actually somewhat simpler, but more general (than a metric that presumes a certain model).

We denote the number of edges in the n th sampled subgraph by the random variable Y_n . For the goodness of fit test, we require its distribution. In the observed network we have $|V|$ nodes which gives $V_e = \binom{|V|}{2}$ possible edges, of which $|E|$ exist. If we have a sampled subgraph of k nodes this has $k_e = \binom{k}{2}$ possible edges. Naively we might estimate the probability of observing y_n : edge in the n th subgraph to be

$$P(Y_n = y_n) = \frac{\binom{|E|}{y} \binom{V_e - |E|}{k_e - y_n}}{\binom{V_e}{k_e}}.$$

This probability comes from the fact that the number of ways of choosing the y_n edges from the total $|E|$ edges in the observed network is $\binom{|E|}{y}$; we must also have chosen $k_e - y_n$ potential edges that were not edges - there are $\binom{V_e - |E|}{k_e - y_n}$ ways of doing this. Therefore $\binom{|E|}{y} \binom{V_e - |E|}{k_e - y_n}$ is the number of ways we can take a sample of k nodes with y_n edges. The total number of samples of size k nodes that we can take from the observed network is $\binom{V_e}{k_e}$, and hence the probability above. Observant readers will notice that this is the probability mass function of a hypergeometric distribution with a population of size V_e with $|E|$ successes and $V_e - |E|$ failures, from which we take a random sample of k_e without replacement and denote the number of successes in the sample as Y_n .

However, the analysis above assumes that the edges are independent of one another in the sampling process, *i.e.*, the selection of an edge does not change the probability of another edge being selected in the same random sample. This is not the case as we will now illustrate.

Consider the two networks given in Figure 1 denoted Network (a) and Network (b). Both of these networks have 4 nodes and two edges. Consider sampling three nodes from each network. The possible samples are given in Table 1.




Figure 1: Two networks illustrating that sampling from networks does not give a hypergeometric distribution.

The resultant probability mass functions and reference hypergeometric are given in Table 2. The distributions are not equivalent to the hypergeometric and not equivalent to

Nodes sampled	Number of edges	
	Network (a)	Network (b)
$\{1, 2, 3\}$	2	1
$\{1, 2, 4\}$	1	1
$\{1, 3, 4\}$	1	1
$\{2, 3, 4\}$	0	1

Table 1: Edges obtained for each of the possible samples of three nodes from Networks (a) and (b) in Figure 1.

	y		
	0	1	2
Network (a)	1/4	1/2	1/4
Network (b)	0	1	0
Hypergeometric(2,4,3)	1/5	3/5	1/5

Table 2: The probability mass function $P(Y = y)$ for the number of edges in a sampled subgraph of size 3 from the networks in Figure 1. For reference the theoretical probabilities calculated from the equivalent hypergeometric distribution are given.

one another. So where is the hypergeometric distribution? Consider all possible networks of four nodes with two edges given in Figure 2. These fifteen networks are all isomorphic to Network (a) or Network (b). Assuming that we know we have a network of four nodes and two edges and sample three nodes, the probabilities $P(Y = y \mid \text{Network } X)$ are given in Table 2, and we can then use the Law of Total Probability to show that $P(Y = y)$ is the same as that given for the hypergeometric distribution given in Table 1. Thus the discrepancy arises because we sample from a single *real* graph, not the complete ensemble of possible random graphs.

As will be seen later, this discrepancy between the true distribution of the number of edges in a subgraph and the hypergeometric distribution becomes negligible as the number of nodes in the observed network becomes large. This means that for large networks, we can use the equivalent hypergeometric distribution to calculate the expected number of edges for the goodness of fit test for homogeneity. The process is outlined in Algorithm 1. We refer to this method as the *Approximation test*.

For small networks, we modify Algorithm 1 so that instead of calculating the P-value using a χ^2 distribution as given in Step 8, we simulate an empirical distribution of X^2 . This is achieved by simulating R GER networks with the same number of nodes and edges as the observed network. For each of these simulated networks, we use Algorithm 1 to calculate $x_r^2, r = 1, \dots, R$, where x_r^2 is the observed value of X^2 (Step 7 of Algo. 1) of the r th network. The empirical P-value, P-value_{emp} is then defined




Figure 2: All possible networks of four nodes with two edges. They are all isomorphic to the networks in Figure 1: twelve to Network (a), and three to Network (b).

Algorithm 1: Goodness of fit test for homogeneity in large networks

Input: Observed network $G(V, E)$

- 1 **for** $n = 1, \dots, N$ **do**
- 2 Sample a subgraph with k nodes from the network G .
- 3 Count the number of edges: y_n
- 4 **end**
- 5 Tabulate y_n into bins according to Algorithm 2 to give the number of subgraphs in each bin: $f_m, m = 1, \dots, M$.
- 6 Calculate the expected number in each bin from the equivalent hypergeometric: $e_m, m = 1, \dots, M$.
- 7 Calculate the χ^2 statistic:

$$X^2 = \sum_{m=1}^M \frac{(f_m - e_m)^2}{e_m}.$$

- 8 Calculate the P-value:

$$P(X^2 \geq x^2), \text{ where } X^2 \sim \chi_{M-1}^2.$$

- 9 **return** $P\text{-value}$
-

$$\text{P-value}_{emp} = \frac{1}{R} \sum_{r=1}^R I(x_r^2 \geq x_{obs}^2),$$

where x_{obs}^2 is the observed value of X^2 for the observed network, and $I(\cdot)$ is an indicator function. This modification we will refer to as the *Empirical Test*. Its advantage is that it is more accurate for small samples, but as we will show it is much less computationally efficient.

4 Results

4.1 Significance levels

We examined the ability of the proposed methods – the Approximation Test and the Empirical Test – to identify networks with heterogeneity by assessing the significance level and power for both methods. To test the significance level, we simulated GER networks of size $|V| = 10^i, i = \{2.00, 2.25, \dots, 4.00\}$ with average node degree, denoted \bar{d} , of 1, 3, 5, and 10 (sparse graphs are more realistic for many applications, but also sparsity makes the estimation problem harder, so we test in this domain). For each of these pairs of parameters $(|V|, \bar{d})$ we simulate 500 GER networks, and for each of these networks we sampled 1000 subgraphs of size k .

The size k was chosen to maximise the variance of the number of edges for the i th subgraph, Y_n , under the assumption that Y_n has a hypergeometric distribution. It can be shown that this is achieved by

$$k = \frac{1 + \sqrt{1 + 2 \times |V|(|V| - 1)}}{2}.$$

We assessed the Empirical Test using the same parameter settings, with the exception that we considered small networks ($|V| \leq 10^3$) only, due to the much larger computation times for this method (see Section 4.4 for details). We also reduced the number of replications from 500 to 200 as a result of the larger computation times. We use 200 samples to simulate the empirical distribution of X^2 in this test: this proved to be a reasonable compromise between accuracy and computation time.

4.2 Significance

For each parameter pair $(|V|, \bar{d})$ we calculated an estimated significance level by calculating the proportion of the P-values and P-value_{emp} that were less than or equal to 0.05. These are shown in Figure 3 for $\bar{d} = 5$ with 95% confidence intervals for the true significance level. The estimated significance levels for the Empirical Test are not

significantly different to 0.05, as desired. On the other hand, the significance level for the Approximation Test were significantly higher than expected for small networks. For $|V| < 1000$ we see values as large as 0.25, indicating a false positive rate of 25%. The significance levels drop back to the correct range at around $|V| = 1000$.

Similar results were obtained for the other values of \bar{d} , as can be see in Table 3. The net result is that one should apply the Approximation Test only for larger networks ($|V| \geq 1000$), but in that range the approximation works quite well.

Approximation Test				Empirical Test			
\bar{d}	$ V $			\bar{d}	$ V $		
	100	1000	10000		100	316	1000
1	0.25	0.08	0.06	1	0.05	0.03	0.07
3	0.23	0.06	0.08	3	0.06	0.04	0.07
5	0.23	0.07	0.05	5	0.07	0.07	0.04
10	0.20	0.06	0.06	10	0.06	0.07	0.03

Table 3: Estimated significance levels obtained for a homogenous network.

4.3 Power

Next we consider the *power* of the methods, by simulating networks that have controlled heterogeneity in the p_{ij} . We considered a modification of the GER network which we call a two-colour GER network (a modification of the models described in Söderberg [2003]). We modify the network to have two type of nodes: denoted as red nodes and blue nodes. There will be $|V_1|$ red nodes and $|V_2|$ blue nodes with $|V_1|$ not necessarily equal to $|V_2|$. The total number of nodes in the network is $|V| = |V_1| + |V_2|$. We also define, for $i > j$

$$p_{ij} = P\{A_{i,j} = 1\} = \begin{cases} p, & i, j \in V_1, \\ q, & i, j \in V_2, \\ \sqrt{pq}, & i \in V_1, j \in V_2, \\ \sqrt{pq}, & i \in V_2, j \in V_1. \end{cases}$$

The network is undirected so $A_{ij} = A_{ji}$ for $i < j$, and as always $A_{ii} = 0$.

The choice of \sqrt{pq} for the edges connecting nodes of different classes was based on the idea that each node contributed to an edge independently. To be consistent with $p_{ij} = p$ for $i, j \in V_1$, a red node should contribute \sqrt{p} , and similarly a blue node \sqrt{q} , and hence $p_{ij} = \sqrt{pq}$ for $i \in V_1, j \in V_2$.

We examine power using two-colour GER networks with larger networks (where the Approximation Test is valid) and average node degrees as considered before. In our case, we set $|V_1| = |V_2| = |V|/2$, *i.e.*, equal group sizes.




Figure 3: Estimated significance level with 95% confidence intervals for the true significance level for $\bar{d} = 5$. Note that the Empirical Test's significance is within bounds of the desire level for all $|V|$, but the Approximation Test only approaches the correct significance near $|V| = 600$.




Figure 4: Computation times (seconds). We can see that the Approximate Test is better than two orders of magnitude faster than the Empirical Test, but that they have the same asymptotic performance in both cases $O(|V|^2)$

We control the degree of heterogeneity through the ratio

$$r = \frac{q - p}{p + q},$$

which is the relative average difference between the two classes' probabilities. We test $r = \{0.01, 0.1, 0.2, 0.5, 0.75, 1\}$. The ratio r is the ratio of the distance of p and q from the average of p and q to the average of p and q . The case $r = 0$ corresponds to homogeneity, and values of r greater than one would result in either p or q being negative and hence are not possible.

For each parameter combination $(|V|, \bar{d}, r)$, we simulated 500 networks. For each we then calculated the P-value using Algorithm 1. We have only considered the power for the Approximate Test due to the excessive computation time of the Empirical Test.

The results are given in Figure 5. As the ratio, r , increases we observe a monotonically increasing power, *i.e.*, the probability of concluding that the two-colour GER network is inhomogeneous increases. It converges to 1, as expected, because as the ratio increases the difference between p and q for a given average value of p and q increases and so we have more heterogenous probabilities for the edges and so should be more likely to conclude that the network is heterogenous.

The power of the test might not appear large for moderately small values of r , but note that given the data, the test is actually very sensitive. For example consider the case where $\bar{d} = 5$, $|V| = 1000$ and $r = 0.5$. In this case, we will pick up 84% of networks with $p - q = 0.005364$: a very small difference in probabilities. We chose to display the results using r as the x -axis in order to make clear comparisons, but that choice hides the difficulty of the probabilities being tested.

As the average node degree \bar{d} increases the estimated power goes to 1 more quickly. This makes sense; networks with larger \bar{d} have on average more edges hence more information. For even moderately non-sparse networks, the test is quite capable.

On the other hand, for a given ratio and \bar{d} , we observe that as $|V|$ increases, the power decreases. This is because larger networks with a given fixed average degree have smaller p and q values, and hence a smaller difference in probabilities (for a given value of r). The smaller difference is slightly harder to detect, even though there are more data points, but note that by the time average degree $\bar{d} \approx 10$ the effect of network size is almost negligible.

4.4 Computational time

We examine the computation times for the two methods on the test networks described in Section 4.2. We show here the results for networks with average node degree $\bar{d} = 5$. We timed each method 100 times on an iMac 2.7GHz Intel Core i5 with 4 cores and 16 GB of RAM.




Figure 5: Estimated power for given two-colour GER networks with ratio r faceted by average node degree for given network size $|V|$. Note that even for moderate r values, the difference in probabilities being detected can be *very small*, e.g., for $\bar{d} = 5$, $|V| = 1000$ and $r = 0.5$, then $p - q = 0.005364$.

The mean computation times for each method and network size is given in Table 4 and Figure 4 from which we see Empirical Test is orders of magnitude slower than Approximate Test, but that both methods are $O(|V|^2)$.

Method	$ V $		
	100	1000	10000
Approximate Test	0.04s	0.05s	0.27s
Empirical Test	7.39s	10.92s	58.23s

Table 4: Average computation times (seconds) for the two methods for calculating P-values for networks with $|V|$ nodes, and $\bar{d} = 5$.

We examine the effect of average node degree on the computation time in Table 5 and Figure 6. Again, we can see that as expected the Empirical Test is more than two orders of magnitude slower than the Approximate Test. The computational cost increases with \bar{d} as $O(|E|)$.

Method	\bar{d}			
	1	3	5	10
Approximate Test	0.04	0.05	0.05	0.06
Empirical Test	8.78	9.74	10.59	13.04

Table 5: Average computation times (seconds) for the two methods for networks with average node degree \bar{d} and $|V| = 1000$.




Figure 6: Computation times (seconds) for each method for network of given average node degree \bar{d} and $|V| = 1000$. We can observe that the times are $O(|E|)$.

5 Application to Australian Research Council data

We will now apply the homogeneity goodness-of-fit test to a real dataset obtained from the Australian Research Council³. This data consists of the Field of Research (FoR) codes entered by applications for all Discovery Project grants applied for between 2010 and 2014 inclusively.

The FoR codes are six-digit number that indicates the areas of research covered by the grant, as outlined by the Australian Bureau of Statistics (ABS)⁴. The codes are hierarchical: that is they consists of three nested 2-digit codes, for example, 010406 is Stochastic Analysis and Modelling which is nested within 0104 (Statistics) which is contained within 01 (Mathematical Sciences). There are 1238 possible 6-digit FoR codes, though only 1174 were used in the dataset.

Each of the 18,476 grant applications in the dataset nominated one or more FoR codes. These are used to help selection of reviewers, and for statistical purposes (for instance, to report the number of grants accepted per research area).

Many grant applications nominate more than one FoR code (the largest number in a single grant was 11). We construct an inter-FoR-code network by creating a link between two FoR codes if there was at least one application that contained both of these FoR codes. Figure 7 shows the 2-digit graph as an illustration of the data, though we analyse the 6-digit graph here. This gave us a network with 1238 nodes and 15,747 edges.




Figure 7: ARC 2-digit FoR-code graph: shows cross-collaborations between 2-digit FoR groupings. Note that at this level the graph is almost complete. We actually analyse the 6-digit graph which is much sparser, but it has too many nodes to usefully visualise. We use the width of the lines to illustrate the number of common grant applications.

³<http://www.arc.gov.au/>

⁴<http://goo.gl/hrWMUh>

We have enough nodes (*i.e.*, $|V| = 1,238 \geq 1000$) to use the Approximation Test to assess homogeneity. We sampled 1000 subgraphs of size 876 nodes. This gave an observed X^2 test statistic of 21,226.12 with a P-value of 0. Therefore, we conclude that we have very strong evidence that there is heterogeneity in the 6-digit FoR-code network.

This is not at all surprising: intuition suggests that cross-collaboration between different fields is dependent on the type of field under consideration. Certain fields are often used within others, for instance Statistics is an important component of many other areas of Science, and hence there is likely to be many bridges between Statistics and Sciences. On the other hand, subjects such as Pure Mathematics exist in relative isolation as a deliberate decision about the nature of the subject. However, while the finding is not surprising, it is reassuring to have a quantitative test, rather than relying on intuition.

6 Discussion

In the beginning, we discussed a problem in network analysis, *i.e.*, that we often have a single observed network (which is a single high-dimensional data-point) from which we would like to choose an appropriate model, and then estimate the parameters for the model. This is often achieved by implicitly assuming homogeneity of the subgraph structure of the network so that subsamples can be exploited to provide multiple data points in model selection or parameter estimation.

We have introduced a simple example of proof of concept of a framework to access this assumption. However our framework can be generalised to elucidate other notions of heterogeneity in the subgraphs.

Our example of using this framework was the concept of homogeneity/heterogeneity in edge probability. We used the number of edges as a simple summary statistic of subsampled graphs. In this example, we have obtained excellent results.

Testing that a network in GER is not that remarkable. It is more the philosophy of the approach – we have proposed a framework to separate local structure from global structure. That is, we are not aiming to test if the network is a GER random graph, but rather we aim to test a property of the network that would be needed if we were to fit the GER random graph to the data.

So how can this framework be applied to other forms of heterogeneity? The outline of the approach is as follows:

- Identify the property that is being explored (in our case edge connection probability), and identify an appropriate model with this property: in our case the GER random graphs.
- Identify summary statistics for the model parameter. Again in our case, we choose

the number of edges in the subgraph to give the most information about the parameter p_{ij} .

- Sample subgraphs from the network and for each one record the summary statistic.
- Perform a goodness of fit test for the summary statistic compared to the theoretical distribution of the summary statistic. The theoretical distribution can be derived from first principles. Or instead, you can derive an empirical distribution for the summary statistic under the assumption of homogeneity by simulating the appropriate network and the sampling from this.

Given this framework, we can now test for heterogeneity of the local network structure, and assess whether a single homogenous model can be applied to the total network, though choosing appropriate sampling and summary statistics might require some creativity for more complex inferences.

7 Conclusion

In this paper, we have illustrated how a sampling procedure can be used to test for heterogeneity in the local structure of a network. We used the simple example of heterogeneity in edge probability of an GER network as a proof of concept, but the method can be applied to other heterogeneities.

We have shown how a goodness of fit can be utilised to test for heterogeneity. In our example we use the number of edges in the subgraphs as the summary statistic and for this we could calculate an approximate distribution using the hypergeometric. An appropriate distribution may not always be available and hence we also shown how an empirical distribution can also be used.

This framework gives the researcher a new method to elucidate heterogeneity in local structure.

An area of future research that we are investigating is using this approach to not only indicate heterogeneity, but to identify its relationship to network structure: a residual type object for network modelling.

Acknowledgements

This work was partially supported by the Australian Research Council through grant DP110103505, and the ARC Centre of Excellence for Mathematical & Statistical Frontiers.

We would also like to thank Aidan Byrne and the ARC for the provision of the For-code dataset used here.

References

- A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286: 509–512, October 1999.
- George E. P. Box and Norman R. Draper. *Response Surfaces, Mixtures and Ridge Analysis*. Wiley, 2nd edition, 2007.
- Jennifer A. Dunne, Richard J. Williams, and Neo D. Martinez. Food-web structure and network theory: The role of connectance and size. *Proceedings of the National Academy of Sciences*, 99(20):12917–12922, 2002. doi: 10.1073/pnas.192407699. URL <http://www.pnas.org/content/99/20/12917.abstract>.
- John P. Huelsenbeck and Fredrik Ronquist. MRBAYES: bayesian inference of phylogenetic trees. *BIOINFORMATICS APPLICATIONS NOTE*, 17(8):754–755, 2001.
- A W Kimball. Errors of the third kind in statistical consulting. *Journal of the American Statistical Association*, 52(278):133, 1957.
- Sang Hoon Lee, Pan-Jun Kim, and Hawoong Jeong. Statistical properties of sampled networks. *Phys. Rev. E*, 73:016102, Jan 2006. doi: 10.1103/PhysRevE.73.016102. URL <http://link.aps.org/doi/10.1103/PhysRevE.73.016102>.
- M. E. J. Newman. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences*, 98(2):404–409, 2001. doi: 10.1073/pnas.98.2.404. URL <http://www.pnas.org/content/98/2/404.abstract>.
- Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 10 lessons from 10 years of measuring and modeling the internet’s autonomous systems. *IEEE Journal on Selected Areas in Communications*, 29:1810–1821, 2011.
- B. Söderberg. Random Graph Models with Hidden Color. *Acta Physica Polonica B*, 34: 5085, October 2003.
- Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.

A Algorithm to tabulate observed number of edges

Algorithm 2: Method to get bins with expected value of 5 or greater according to $hyper(m, n, k)$

input : Parameters for hypergeometric distribution to test for:

m number of successes in population.

n number of failures in population.

k sample size.

N Number of observations

```
1  $c \leftarrow 5/N$ 
2  $p \leftarrow c$ 
3  $j \leftarrow 1$ 
4 while  $p + c < 1 - c$  do
5    $x_j \leftarrow P(X \leq x_j) = p$  where  $X \sim hyper(m, n, k)$ 
6    $p \leftarrow P(X \leq x_j) + c$ 
7    $j \leftarrow j + 1$ 
8 end
9 return  $\{x_1, \dots\}$ 
```
